Advances in Water Resources 94 (2016) 498–509

Contents lists available at ScienceDirect

Advances in Water Resources

journal homepage: www.elsevier.com/locate/advwatres

Multilevel Monte Carlo methods for computing failure probability of porous media flow systems

F. Fagerlund[a], F. Hellman[b,*], A. Målqvist[c], A. Niemi[a]

[a] Department of Earth Sciences, Uppsala University, Box 337, SE-751 05 Uppsala, Sweden
[b] Department of Information Technology, Uppsala University, Box 337, SE-751 05 Uppsala, Sweden
[c] Department of Mathematical Sciences, Chalmers University of Technology and University of Gothenburg, SE-412 96 Göteborg, Sweden

ARTICLE INFO

Article history:
Received 13 November 2015
Revised 26 May 2016
Accepted 10 June 2016
Available online 15 June 2016

Keywords:
CDF estimation
Failure probability
Porous media flow simulation
Multilevel Monte Carlo
Selective refinement

ABSTRACT

We study improvements of the standard and multilevel Monte Carlo method for point evaluation of the cumulative distribution function (failure probability) applied to porous media two-phase flow simulations with uncertain permeability. To illustrate the methods, we study an injection scenario where we consider sweep efficiency of the injected phase as quantity of interest and seek the probability that this quantity of interest is smaller than a critical value. In the sampling procedure, we use computable error bounds on the sweep efficiency functional to identify small subsets of realizations to solve highest accuracy by means of what we call selective refinement. We quantify the performance gains possible by using selective refinement in combination with both the standard and multilevel Monte Carlo method. We also identify issues in the process of practical implementation of the methods. We conclude that significant savings in computational cost are possible for failure probability estimation in a realistic setting using the selective refinement technique, both in combination with standard and multilevel Monte Carlo.

1. Introduction

Simulation of fluid flow in geological formations of porous and fractured media is a tool used to help understand the processes governing groundwater flow, enhanced oil recovery and underground carbon storage. Simulations can help assess the suitability for enhanced oil recovery or carbon storage at potential injection sites. There are, however, typically large uncertainties in the physical properties (e.g. permeability, porosity, caprock topology) of the porous medium. In this work, we consider forward propagation of uncertainties in input data through the governing flow model to a scalar quantity of interest. The deterministic problems are typically computationally expensive to solve. Adding a high-dimensional stochastic input space poses an even greater challenge. It is not always necessary to compute the full distribution of the quantity of interest, but it is common to seek, for example, the mean value or variance of the quantity of interest. Here, we seek the probability for the quantity of interest to be less than a fixed critical value, an event that we call a failure. We assume that the random input data follows a known distribution.

Although the methods presented in this paper are applicable to more general situations, this work focuses on a specific, but common problem for two-phase fluid flow in porous media. We consider a scenario where a non-wetting fluid is injected into a reservoir initially filled with a wetting fluid. We compute the reservoir volume swept by the injected fluid after a certain time. This is a measure of the sweep efficiency. One situation where this problem is of interest is when estimating the capacity of a geological formation to store carbon dioxide. In this context we consider it a failure of the carbon dioxide storage system if the estimated storage capacity of a target formation is lower than some critical value needed for the planned storage operations. Sweep efficiency is one important indicator for storage capacity (Lake, 1989; Shafeen et al., 2004; Tanaka et al., 1995). We consider sweep efficiency as our quantity of interest and the permeability field as a random input data. We then seek to estimate the probability that the sweep efficiency is smaller than some critical value. We call this the failure probability. We use a macroscopic incompressible two-phase fractional flow model based on Darcy's law and conservation of mass, which we solve using standard finite element and finite volume techniques.

The general setting is that we have a quantity of interest X defined on random input sample space Ω and that we seek $p = \Pr(X \leq y)$ for some critical value y. For our problem, X is the sweep efficiency of the injected phase. The sample space Ω contains the

* Corresponding author.
E-mail addresses: fritjof.fagerlund@geo.uu.se (F. Fagerlund), fredrik.hellman@it.uu.se (F. Hellman), axel@chalmers.se (A. Målqvist), auli.niemi@geo.uu.se (A. Niemi).

http://dx.doi.org/10.1016/j.advwatres.2016.06.007
0309-1708/

possible permeability fields. The failure probability p is the value of the cumulative distribution function (cdf) for the sweep efficiency functional X at the critical value y. In this work, we study how error bounds on this functional can be used to improve the performance of standard and multilevel Monte Carlo methods. This is a continuation of the papers (Elfverson et al., 2014, 2016) where the selective refinement technique was introduced. While those two papers focus entirely on asymptotic cost rate results (for large sample sizes), this work applies the methods presented therein to more realistic problems in order to investigate to what extent the predicted gains in computational efficiency can be realized in practice. In addition, special care is taken to describe the construction of the error bounds necessary to use the selective refinement approach. In total four methods of Monte Carlo type are studied: Standard Monte Carlo, multilevel Monte Carlo (Giles, 2008) and the two previous Monte Carlo methods in combination with selective refinement (Elfverson et al., 2016). We briefly review the methods and related works below.

An estimate of the failure probability in a Monte Carlo (MC) estimator is obtained by computing the mean of a set of zeros and ones, from independent simulations yielding successful or failing scenarios, respectively. An improvement of the MC estimator for application on numerical simulations of controllable numerical accuracy is the multilevel Monte Carlo (MLMC) estimator (Heinrich, 2001) first introduced in the context of differential equations in Giles (2008). It has since then been applied to elliptic PDEs in Barth et al. (2011); Cliffe et al. (2011); Park and Teckentrup (2015); Teckentrup et al. (2013) and has been further analyzed and extended in Collier et al. (2015); Haji-Ali et al. (2015a,b). It exploits the convergence of numerical solutions with respect to some discretization parameter h (typically mesh size) and uses a series of corrector estimators of increasing cost but decreasing variance to allow for redistribution of the variance reduction effort to cheap low-accuracy problems. MLMC methods for point evaluation of a cdf is also studied in Avikainen (2009); Giles et al. (2015). In Avikainen (2009) they investigate how the convergence rate of the error in irregular functionals of the quantity of interest approximation (in particular the failure probability functional, or binary option payoff) depends on the convergence rate of the error in the quantity of interest approximation itself. This motivates that the irregular failure probability functional can be used in a multilevel Monte Carlo method. In Giles et al. (2015), the failure probability functional is smoothed to improve the convergence rates. In contrast, the methods considered in the present work are all based on the irregular functional, i.e. no smoothing is used. Another (independent) improvement is Monte Carlo with selective refinement (MC-SR, Elfverson et al. (2014)) for estimation of p-quantiles or point evaluation of a cdf. Selective refinement uses error bounds of the quantity of interest to determine which realizations need to be solved on a fine mesh, and which realizations can be solved on a coarser mesh to a smaller cost. Selective refinement can be combined with multilevel Monte Carlo (MLMC-SR, Elfverson et al. (2016)). In this paper, the computational cost of the four setups (MC, MC-SR, MLMC and MLMC-SR) are estimated for the sweep efficiency in a two-phase flow scenario where the failure probability is of magnitude 5–10% and an absolute accuracy of this probability in the order of a few percent is required. This work is to a large extent experimental and also aims at identifying problems and difficulties in the practical implementation of the mentioned improved Monte Carlo techniques. In particular, the two issues of estimating the variance of the correctors for the multilevel Monte Carlo method, and the establishment of an error bound required for selective refinement are addressed.

The paper is structured as follows. The problem setting and the continuous two-phase flow model is described in Section 2. In Section 3 we introduce a mesh hierarchy, the space and time discretization used, and the procedure for generating random permeability fields. Section 4 gives an overview of the four Monte Carlo setups and their asymptotic computational cost rates. Two preparatory experiments (specific to the sweep efficiency problem) where a computable error bound, and variances and cost models are estimated are presented in Section 5. In Section 6 we describe the experiment that quantifies the computational cost for the four methods. We conclude with a summary and discussion in Section 7.

2. Continuous model

The problem is to estimate the failure probability $p = F(y) = \Pr(X \leq y)$, where y is a given value, called critical value, and F is the cdf of the sweep efficiency X. The sweep efficiency is modeled as a functional of the solution to a nonlinear PDE with random inputs modeling two-phase flow. This section describes the continuous model for the two-phase flow system, introducing the PDE in Section 2.1, random input in Section 2.2 and sweep efficiency in Section 2.3. We emphasize that the model presented here should be considered an example. The methods presented in Section 4 are applicable in a much more general setting.

2.1. Fractional-flow formulation for two-phase flow

We use the fractional flow equations as model for the two-phase flow in porous media and assume isotropic permeability, immiscibility, incompressibility, no capillary forces and that the flow is perpendicular to the gravitational field. Let the domain be the two-dimensional unit square and denoted by $\mathcal{D} = [0,1]^2$ and its boundary by Γ. We denote an arbitrary phase by α and the two particular phases by $\alpha = w$ and $\alpha = n$ for the wetting and non-wetting phase, respectively. For each phase, we have a mass conservation equation

$$\rho_\alpha \phi \frac{\partial s_\alpha}{\partial t} + \rho_\alpha \nabla \cdot \mathbf{u}_\alpha = \nu_\alpha, \quad \alpha = w, n, \tag{1}$$

in $\mathcal{D}$, where ρ_α is density, ϕ is porosity, s_α is saturation, $\mathbf{u}_\alpha$ is volumetric flux and ν_α is a source term. The pore space is occupied only by the two fluids, i.e. $s_w + s_n = 1$. For convenience, we denote the wetting phase saturation by $s = s_w = 1 - s_n$ and refer to it simply by saturation. The flux is coupled with pressure and saturation via the relative permeabilities in Darcy's law,

$$\mathbf{u}_\alpha = -\frac{k_{r,\alpha}(s)k}{\mu_\alpha}\nabla p, \quad \alpha = w, n. \tag{2}$$

Here, k is the isotropic permeability field, $k_{r,\alpha}$ is the relative permeability, p is pressure (assumed equal for both phases), and μ_α is the dynamic viscosity. For the relative permeability, we use the following power law (a modification of the Brooks–Corey relation (Brooks and Corey, 1964))

$$k_{r,w} = (s_e)^3 \quad \text{and} \quad k_{r,n} = \zeta(1 - s_e)^3 \tag{3}$$

for the wetting and non-wetting phase, respectively, where s_e is the effective wetting fluid saturation $s_e = (s_w - s_{r,w})(1 - s_{r,w} - s_{r,n})^{-1}$. Here, $s_{r,\alpha}$ is the residual saturation for the two phases and ζ is a parameter, which we set to 1.

We now present the fractional flow formulation. We denote the total fluid flux by $\mathbf{u} = \mathbf{u}_w + \mathbf{u}_n$ and the phase mobilities for the two phases by

$$\lambda_\alpha(s) = \frac{k_{r,\alpha}(s)}{\mu_\alpha}, \quad \alpha = w, n. \tag{4}$$

The total mobility is defined as $\lambda(s) = \lambda_w(s) + \lambda_n(s)$ and the fractional flow function as $f(s) = \lambda(s)^{-1}\lambda_w(s)$. The wetting phase fluxes can be expressed in terms of total flux using the fractional